arXiv:1606.09383v1 [cs.LG] 30 Jun 2016

On Approximate Dynamic Programming with Multivariate Splines for Adaptive Control

W.J. Eerland*, C.C. de Visser†, E. van Kampen‡

Abstract

We define a SDP framework based on the RLS TD algorithm and multivariate simplex B-splines. We introduce a local forget factor capable of preserving the continuity of the simplex splines. This local forget factor is integrated with the RLS TD algorithm, resulting in a modified RLS TD algorithm that is capable of tracking time-varying systems. We present the results of two numerical experiments, one validating SDP and comparing it with NDP and another to show the advantages of the modified RLS TD algorithm over the original. While SDP requires more computations per time-step, the experiment shows that for the same amount of function approximator parameters, there is an increase in performance in terms of stability and learning rate compared to NDP. The second experiment shows that SDP in combination with the modified RLS TD algorithm allows for faster recovery compared to the original RLS TD algorithm when system parameters are altered, paving the way for an adaptive high-performance non-linear control method.

1 Introduction

Bellman described multi-stage decision processes from a mathematical point of view in Bellman [1957], this algorithm was called Dynamic Programming (DP). The DP principle comes down to describing each state with a value in a value function and moving the system to the state with the highest value. The value function is also called the cost-to-go function and stores the expected sum of future rewards for each state. In general this function cannot be found directly, therefore an iterative approach like Temporal Difference (TD)-learning [Sutton and Barto, 1998, Sutton, 1988] is used.

DP has the ability to solve complex control problems in diverse environments, and it is possible to view the environment as a black-box by modelling it using system identification techniques. The advantage of combining DP and system identification techniques, is that it leads to an adaptive control scheme. These adaptive controllers have already been successfully trained off-line for many purposes, ranging from agile missile interception [Han and Balakrishnan, 1999] to aircraft auto-landing and control [Saini and Balakrishnan, 1997]. Also,

*University of Southampton, Transportation Research Group [w.j.eerland@soton.ac.uk]
†Delft University of Technology, Control and Simulation
‡Delft University of Technology, Control and Simulation

an on-line adaptive critic flight control was implemented on a six-degree-of-freedom business jet aircraft over its full operating envelope, improving its performance when unexpected conditions are encountered for the first time [Ferrari and Stengel, 2004].

Barto, Sutton and Anderson used neural networks to parametrise the value function [Barto et al., 1983]. However, this function approximator is non-linear in the parameters, with the result that stability can only be guaranteed if bounded network weights are used, where bounds are determined by off-line analysis. This approach has been applied to examples [Han and Balakrishnan, 1999, Saini and Balakrishnan, 1997, Ferrari and Stengel, 2004] mentioned before. More recently, an adaptive controller has been introduced in Anderson et al. [2007], and again the neural network weights are bounded to guarantee stability. A drawback is that for a time-varying system these bounds shift and stability can no longer be guaranteed. This combination of neural networks and DP is commonly referred to as Neuro Dynamic Programming (NDP) [Bertsekas and Tsitsiklis, 1996].

There is a proof of convergence when linear-in-the-parameters function approximators are used in Tsitsiklis and Van Roy [1997], however, this proof demands knowledge of the shape of the optimal value function. With the Recursive Least Squares Temporal Difference (RLS TD) algorithm convergence in a stochastic framework is assured [Bradtke and Barto, 1996], even when the linear regression basis cannot perfectly fit the value function. In the last decade, the DP theory in continuous time and space has been further developed in Doya [1996, 2000], Morimoto and Doya [2005]. More recently, the proof of convergence has been extended to include the optimal policy [Ma and Powell, 2009], however one of the problems that remains is the a-priori unknown shape of the value function.

Using polynomials as a function approximator in a DP framework was investigated by Bellman in 1963 [Bellman et al., 1963]. In 1985, Schweitzer [1985] concentrates on the use of global polynomials in combination with DP. With the development of the RLS TD algorithm, it is possible to obtain a proven convergence by combining it with a polynomial approximation as discussed in Ma and Powell [2009]. However the limitation is that the approximation power of global polynomials can only be increased by increasing the order of the polynomials, which will also lead to numerical instabilities in the solution schemes of the approximation. According to Summers et al. [2013], using the sum of squares allows the use of higher order polynomials, but will eventually still lead to numerical instability.

A recommendation in Ma and Powell [2009] is to use local polynomial regression to treat Markov Decision Process (MDP) problems with value functions of unknown form. This recommendation is supported by Daniel [1976], which states that it is highly desirable from an efficiency point of view to use local polynomial regression. Recently a novel method based on multivariate simplex B-splines has been applied in a linear regression framework [de Visser et al., 2009]. The use of a local polynomial basis allows for transparency and efficient (sparse) computational schemes. Furthermore, the spatial location of the B-coefficients and modularity of the triangulation allow for local model modification and refinement [Lai and Schumaker, 2007]. These properties make multivariate simplex B-splines an excellent candidate for use in the DP framework.

There are existing approaches which combines Multivariate Adaptive Regression Splines (MARS) [Friedman, 1991] and DP [Chen, 1999, Chen et al., 1999, Cervellera et al., 2007]. However, multivariate simplex B-splines distinguishes itself from MARS in terms of computational efficiency by using B-splines [Bakin et al., 2000]. And they are supported by a triangulation of simplices, allowing functionality in a non-square domain [Lai and Schumaker, 2007].

The contribution of this paper is a framework that allows the use of multivariate simplex B-splines in combination with the RLS TD algorithm, giving rise to Spline Dynamic Programming (SDP) that enables control of non-linear stochastic systems. A method for continuous local value function adaptation is presented which is enabled by the spatial location property of the coefficients of the multivariate splines; this is achieved by implementing a new formulation for the covariance update step. The effectiveness of this SDP framework is investigated by comparing it with NDP in terms of computational complexity and performance. Furthermore, the RLS TD algorithm is modified to allow for adaptive control of time-varying systems. The validation and comparison of both cases are investigated with a non-linear 2D control problem, the pendulum swing-up.

In section 2, we briefly introduce the DP framework and show how the value function and greedy policy are constructed. In section 3, we give a brief introduction on the mathematical background of multivariate simplex B-splines, the function approximator used in the SDP framework. The SDP framework itself is explained in section 4; the main purpose of this section is to address the steps specific to using multivariate simplex B-splines for value function approximation in combination with the RLS TD algorithm. Section 5 is there to demonstrate that the SDP framework indeed works for the given control problem. SDP is compared with neural networks for system performance on a stochastic system and a time-varying system. These results are discussed in section 6 and finally conclusions and recommendations are presented in section 7.

2 Preliminaries on Dynamic Programming

In this section, we present the preliminaries on DP, the algorithm that is part of the SDP framework. For a more complete description we refer to Sutton and Barto [1998], Si et al. [2004], Powell [2007], Busoniu et al. [2010], Bertsekas [2007]. We start with a brief overview of the MDP followed by the policy evaluation and concluded with the policy improvement.

2.1 Markov Decision Processes

The policy evaluation problem associated with discrete-time stochastic optimal control problems is referred to as a MDP. Finding the solution to an MDP is a sequential optimization problem where the goal is to find a policy that maximizes the sum of the expected infinite-horizon discounted rewards.

Let $\mathbf{x}_t \in X$ be the state vector and $\mathbf{u}_t \in U$ be the input vector, both at time t, where $\mathbf{u}_t$ is determined by the policy π and X and U denote the finite sets of states and inputs. The reward function is $r_{t+1}(\mathbf{x}_{t+1}, \mathbf{u}_t)$ and $0 \leq \gamma < 1$ is the discount factor. The goal is to find a policy π that obtains the maximum total reward.

For each policy π there exists a value function $V^\pi(\mathbf{x}_t)$ that indicates a measure of long-term performance at each state:

$$V^\pi(\mathbf{x}_t) = \sum_{k=t}^{\infty} \gamma^{k-t} r_{k+1}(\mathbf{x}_{k+1}, \mathbf{u}_k) \tag{1}$$

The objective can now be formulated as finding a policy π^* such that $V^{\pi^*}(\mathbf{x}) \geq V^\pi(\mathbf{x})$ for all $\mathbf{x} \in X$ and for all policies π. This policy π^* is called the optimal policy and can be found by applying both policy evaluation and policy improvement Sutton and Barto [1998].

The policy evaluation determines the $V^\pi(\mathbf{x}_t)$ of the current policy π, where the policy improvement uses this knowledge to adjust the policy π such that it ends up in the most valuable states.

2.2 Policy evaluation

In order to evaluate the value function in an iterative fashion, Sutton [1988] uses TD-learning. For TD-learning, the following must hold:

$$\begin{aligned} V^\pi(\mathbf{x}_t) = &\quad r_{t+1} + \textstyle\sum_{k=t+1}^{\infty} \gamma^{k-t} r_{k+1} \\ = &\quad r_{t+1} + \gamma V^\pi(\mathbf{x}_{t+1}) \end{aligned} \tag{2}$$

If this equality does not hold, the difference is called the TD-error:

$$e_t = r_{t+1} + \gamma V^\pi(\mathbf{x}_{t+1}) - V^\pi(\mathbf{x}_t) \tag{3}$$

By minimizing this TD-error, the value function can be constructed in an iterative approach. In order to construct a value function for a continuous problem, there is need for parameterization to describe a complete state space with a finite number of parameters. The value function now becomes $V^\pi(\mathbf{x}_t, \mathbf{c}_t)$, where $\mathbf{c}_t$ are the parameters at time t that shape the continuous value function in domain X.

In Bradtke and Barto [1996] the RLS TD algorithm was introduced; this algorithm and its computational complexity is visible in Table 1. Here $\hat{d}$ and $\hat{a}$ represent the number of coefficients per simplex and total number of coefficients respectively, $\mathbf{B}$ represents the linear regression matrix and $\mathbf{P}$ is the parameter covariance matrix. These parameters will be further explained in section 3. RLS TD converges to the least squares approximation of the optimal value function $\mathbf{c}^*$, given that each state $\mathbf{x} \in X$ is visited infinitely often. Although RLS TD requires more computations per time-step than TD(λ) algorithms [Sutton, 1988], it is more efficient in the statistical sense as more information is extracted from training experience, allowing it to converge faster [Bradtke and Barto, 1996]. Furthermore, while Least Mean Square (LMS) aims to decrease the mean square error e_t at each time-step separately, RLS TD minimizes this objective function:

$$\mathcal{J}_t = \frac{1}{t} \sum_{k=0}^{t} [e_k]^2 \tag{4}$$

Table 1: RLS TD algorithm, from Bradtke and Barto [1996]

Step		Action	Computational Complexity
(1)	$e_t =$	$r_{t+1} - (\mathbf{B}_t - \gamma\mathbf{B}_{t+1})^\top \mathbf{c}_t$	$\mathcal{O}(\hat{d})$
(2)	$\mathbf{P}_{t+1} =$	$\mathbf{P}_t - \frac{\mathbf{P}_t\mathbf{B}_t(\mathbf{B}_t-\gamma\mathbf{B}_{t+1})^\top\mathbf{P}_t}{1+(\mathbf{B}_t-\gamma\mathbf{B}_{t+1})^\top\mathbf{P}_t\mathbf{B}_t}$	$\mathcal{O}(\hat{a}^2)$
(3)	$\mathbf{c}_{t+1} =$	$\mathbf{c}_t + \frac{\mathbf{P}_t}{1+(\mathbf{B}_t-\gamma\mathbf{B}_{t+1})^\top\mathbf{P}_t\mathbf{B}_t}\mathbf{B}_t e_t$	$\mathcal{O}(\hat{a}^2)$

2.3 Policy improvement

The computation of the value function is called policy evaluation. Using this value function, a greedy action can be selected. If a policy is updated in this manner, it is called (greedy) policy improvement:

$$\mathbf{u}_t(\mathbf{x}_t) = \max_{\mathbf{u}_t \in U} [V^\pi(\mathbf{x}_t)] \tag{5}$$

The repetition of the policy evaluation and policy improvement is called policy iteration and will result in an optimal policy [Sutton and Barto, 1998]. According to Doya [1996], the optimal non-linear feedback control law is a function of value function's gradient:

$$\mathbf{u}_t(\mathbf{x}_t) = \mathbf{u}^{\max}\ g\left(\frac{1}{c}\tau\frac{\partial V^\pi(\mathbf{x}_t)}{\partial \mathbf{x}_t}\frac{\partial f(\mathbf{x}_t,\mathbf{u}_t)}{\partial \mathbf{u}_t}\right) \tag{6}$$

where $\mathbf{u}^{\max}$ is the maximum control input, $g(x) = \tanh(\frac{\pi}{2}x)$, c is the control cost parameter, τ is the step-size parameter and $f(\mathbf{x}_t, \mathbf{u}_t)$ are the system dynamics.

This optimal control law is applied to the pendulum swing-up task, with the result visible in A. Having discussed both the policy evaluation and improvement, section 3 will now describe the parametrization of the value function using multivariate simplex B-splines.

3 Preliminaries on Multivariate Simplex B-Splines

This section serves as a brief introduction to the mathematical theory of the simplex B-splines. For a more extensive and general introduction to multivariate spline theory we refer to Lai and Schumaker [2007]. We start by introducing the basic concept of a single basis polynomial and B-form, then introduce the triangulation, followed by the vector notation of the B-form. Finally the Recursive Least Squares (RLS) estimator for simplex splines is reviewed.

3.1 Simplex and barycentric coordinates

The polynomial basis of a multivariate simplex B-spline is defined on a simplex. A simplex is defined by the non-degenerate vertices $(\upsilon_0, \upsilon_1, \ldots, \upsilon_n) \in \mathbb{R}^n$ and thus creates a span in n-dimensional space. Any point $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ can be transformed to a barycentric coordinate $\mathbf{b}(\mathbf{x}) = (b_0, b_1, \ldots, b_n)$ with respect to a simplex. The relation between Cartesian coordinate $\mathbf{x}$ and barycentric coordinate $\mathbf{b}(\mathbf{x})$ is:

$$\mathbf{x} = \sum_{i=0}^{n} b_i \upsilon_i \qquad \sum_{i=0}^{n} b_i = 1 \tag{7}$$

3.2 Triangulation

Any number of simplices can be combined into a triangulation, where a triangulation $\mathcal{T}$ is a special partitioning of a domain into a set of J non-overlapping simplices and is defined in Lai and Schumaker [2007] as:

$$\mathcal{T} \equiv \bigcup_{i=1}^{J} t_i, \quad t_i \cap t_j \in \{\emptyset, \tilde{t}\}, \quad \forall t_i, t_j \in \mathcal{T} \quad , i \neq j \tag{8}$$

with $\tilde{t}$ a simplex of dimension ($< n$). A popular triangulation method is Delaunay triangulation [Lee and Schachter, 1980]. Recently, a new method for creating globally optimal triangulations named Intersplines was introduced in de Visser et al. [2012].

3.3 Basis functions

The polynomial basis of the simplex spline are the Bernstein basis polynomials $B^d_\kappa(\mathbf{b})$, where d is the degree of the spline and $\mathbf{b}$ is the barycentric coordinate discussed earlier:

$$B^d_\kappa(\mathbf{b}) = \frac{d!}{\kappa!}\mathbf{b}^\kappa \tag{9}$$

here $\kappa \geq 0$ is a multi-index, which has properties: $\kappa! = \kappa_0!\kappa_1! \cdots \kappa_n!$, $|\kappa| = \kappa_0 + \kappa_1 + \cdots + \kappa_n$ and $\mathbf{b}^\kappa = b_0^{\kappa_0} b_1^{\kappa_1} \cdots b_n^{\kappa_n}$. The valid permutations of κ with the constraint $|\kappa| = d$ equal the total number of B-coefficients and basis polynomials per simplex and is equal to:

$$\hat{d} = \frac{(d+n)!}{n!d!} \tag{10}$$

In combination with a set of J non-overlapping simplices, the total number of B-coefficients for a complete triangulation is:

$$\hat{a} = J \cdot \hat{d} \tag{11}$$

The multi-index κ has a requirement on the ordering, called a lexicographical sorting order which is introduced in Lai and Schumaker [2007]. This means $\kappa_{\nu\mu\varkappa}$ comes before κ_{ijk} provided that $\nu > i$, or if $\nu = i$, then $\mu > j$, or if $\nu = i$ and $\mu = j$, then $\varkappa > k$. Thus for $d = 2$ the order is:

$$\kappa_{2,0,0}, \quad \kappa_{1,1,0}, \quad \kappa_{1,0,1}, \quad \kappa_{0,2,0}, \quad \kappa_{0,1,1}, \quad \kappa_{0,0,2} \tag{12}$$

Each B-coefficient has a unique position within the simplex based on κ; this relation between B-coefficient and spatial position allows the creation of a B-net as seen in Figure 1.

3.4 Vector formulation of the B-form

In order to complete the vector formulation, $\mathbf{B}^d_{t_j}(\mathbf{b}) \in \mathbb{R}^{\hat{d}\times 1}$ is introduced as a vector of Bernstein basis polynomials (Eq. 9) of simplex t_j, which are sorted lexicographically as indicated by Eq. 12. Adapted from de Visser and Verhaegen [2013], we define the B-form as a row vector on simplex t_j:

$$p(\mathbf{b}) = \begin{cases} \mathbf{B}^d(\mathbf{b})^\top \cdot \mathbf{c}^{t_j}, & \mathbf{x} \in t_j \\ 0, & \mathbf{x} \notin t_j \end{cases} \tag{13}$$

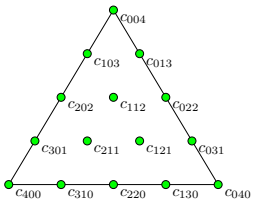

Figure 1: B-net overview of a 4^{th} degree simplex spline

where $\mathbf{c}^{t_j}$ are the B-coefficients on simplex t_j. The matrix operation to evaluate the simplex B-spline function of degree d and continuity order r, defined on a triangulation $\mathcal{T}_J$ is:

$$s_d^r(\mathbf{b}) \equiv \mathbf{B}(\mathbf{b})^\top \cdot \mathbf{c} \in \mathbb{R}, \quad \mathbf{x} \in \mathcal{T}_J \tag{14}$$

Now $\mathbf{B}(\mathbf{b})$ (note the absence of the superscript 'd') is the global vector of basis polynomials:

$$\mathbf{B}(\mathbf{b}) \equiv \left[\mathbf{B}^d(\mathbf{b})_{t_1}^\top \ \mathbf{B}^d(\mathbf{b})_{t_2}^\top \ \cdots \ \mathbf{B}^d(\mathbf{b})_{t_J}^\top\right]^\top \in \mathbb{R}^{\hat{a} \ \mathtt{x} \ 1} \tag{15}$$

The global vector of B-coefficients $\mathbf{c}$ is defined as:

$$\mathbf{c} \equiv \left[\mathbf{c}^{t_1\top} \ \mathbf{c}^{t_2\top} \ \cdots \ \mathbf{c}^{t_J\top}\right]^\top \in \mathbb{R}^{\hat{a} \ \mathtt{x} \ 1} \tag{16}$$

The spline space is the space of all spline functions s_d^r in the triangulation $\mathcal{T}$. We use the definition of the spline space from Lai and Schumaker [2007]:

$$S_d^r(\mathcal{T}) \equiv \{s_d^r \in C^r(\mathcal{T}) : s_d^r|_t \in \mathbb{P}_d, \forall t \in \mathcal{T}\} \tag{17}$$

where $\mathbb{P}_d$ is the space of polynomials of degree d.

3.5 Continuity

Since $p(\mathbf{b})$ is a linear combination of continuous functions, $s_d^r(\mathbf{b})$ is naturally continuous on each simplex. However, in order to assure continuity of $S_d^r(\mathcal{T})$ between simplices, constraints are imposed on the relations between the coefficients of different simplices. The continuity order r fixes the derivatives $\frac{d^r p}{db^r}$ on the edges between neighboring simplices. The required continuity conditions can be calculated using Lai and Schumaker [2007]:

$$\begin{aligned} c^{t_i}_{(\kappa_0,\ldots,\kappa_{n-1},m)} &= \textstyle\sum_{|\gamma|=m} c^{t_j}_{(\kappa_0,\ldots,\kappa_{n-1},0)+\gamma} B^m_\gamma(w) \\ &\qquad\qquad 0 \le m \le r \end{aligned} \tag{18}$$

Here w is a vertex of simplex t_j which is not found on the edge that is shared with simplex t_i. All constraints required for continuity are collected in the smoothness matrix $\mathbf{H}$, with each row containing a new constraint and the columns consisting of the coefficients de Visser et al. [2009, 2011]. These equations are all equaled to zero, resulting in the following matrix form:

$$\mathbf{H}\mathbf{c} = 0 \tag{19}$$

with $\mathbf{c}$ as in Eq. 16.

3.6 Approximation power

To describe the approximation power, the following definition from Lai and Schumaker [2007], chapter 10.1 is used:

Definition *(Approximation power of $S_d^r(\mathcal{T})$)* Fix $0 \leq r < d$ and $0 < \theta \leq \pi/3$. Let m be the largest integer such that for every polygonal domain Ω and every regular triangulation $\mathcal{T}$ of Ω with smallest angle θ, for every $f \in W_q^m(\Omega)$, there exists a spline $s \in S_d^r(\mathcal{T})$ with

$$||f - s||_{q,\Omega} \leq K \ |\mathcal{T}|^m \ |f|_{m,q,\Omega} \tag{20}$$

where the constant K depends only on r, d, θ, and the Lipschitz constant of the boundary of Ω. Then we say that S_d^r has approximation power m in the q-norm. If this holds for $m = d + 1$, we say that S_d^r has full approximation power in the q-norm.

The theory behind this is extensive and available in Lai and Schumaker [2007], however for now it is important to realize that $|\mathcal{T}|^m$ is a function of the longest edge in triangulation $\mathcal{T}$. This reveals that it is possible to locally increase the approximation power by reducing the length of the longest edge in $\mathcal{T}$.

3.7 Recursive least squares

RLS is a method which allows the estimated parameters $\mathbf{c}$ to be updated online with the use of the parameter covariance matrix $\mathbf{P}$ de Visser et al. [2011]. The algorithm and the computational complexity is found in Table 2. Note that it is essential to keep the column relations of $\mathbf{P}$ intact to enforce the constraints $\mathbf{Hc} = 0$ from Eq. 19. To initialize the $\mathbf{P}$ matrix, we use $\mathbf{Z}$, the orthogonal projector on the null-space of $\mathbf{H}$ Lawson and Hanson [1974]:

$$\mathbf{P}_1 = \beta_1 \ \mathbf{Z} \tag{21}$$

where $\mathbf{Z} = (\mathbf{I} - \mathbf{H}^{+}\mathbf{H})$, in which $\mathbf{I}$ and $\mathbf{H}$ are the identity and smoothness matrix respectively. The parameter $\beta_1 > 0$ indicates the confidence in the initial estimated parameters, where a larger β_1 indicates a lower confidence level. To initialize $\mathbf{c}$, it is important to pick these such that the continuity constraints are not violated. For this, a constrained Least Squares (LS) fit of the estimated shape can be used, using the approach from de Visser et al. [2009]. If no knowledge is available, initialization of all coefficients at zero will satisfy the constraints, resulting in:

$$\mathbf{c}_1 = \mathbf{0} \tag{22}$$

4 Spline Dynamic Programming

This section has a dual purpose; it will introduce the framework that combines both simplex splines and the RLS TD algorithm, and the modified RLS TD algorithm with the ability to track time-varying systems will be defined.

Table 2: RLS algorithm from de Visser et al. [2011]

Step		Action	Computational Complexity
(1)	$\epsilon_t =$	$y_t - \mathbf{B}_t^\top \mathbf{c}_t$	$\mathcal{O}(\hat{d})$
(2)	$\mathbf{P}_{t+1} =$	$\mathbf{P}_t - \frac{\mathbf{P}_t \mathbf{B}_t \mathbf{B}_t^\top \mathbf{P}_t}{1+\mathbf{B}_t^\top \mathbf{P}_t \mathbf{B}_t}$	$\mathcal{O}(\hat{a}^2)$
(3)	$\mathbf{c}_{t+1} =$	$\mathbf{c}_t + \mathbf{P}_{t+1} \mathbf{B}_t \epsilon_t$	$\mathcal{O}(\hat{a}^2)$

4.1 The SDP framework

To successfully represent the optimal value function with a simplex spline, a spline space has to have sufficient approximation power at each point of the value function domain. While it is theoretically possible to have an infinite refinement in terms of triangulation, the idea behind the parametrization of the value function is that there is no need for an infinite amount of states, but only a limited amount of parameters to describe the entire state space X. With no a-priori information available, a triangulation consisting of nodes positioned in a grid is a good initial estimate, as it evenly distributes the approximation power over the domain according to Eq. 20. What remains is to construct a spline space $S_d^r(\mathcal{T})$ are the polynomial degree and continuity order, parameters with a global effect on the spline function. Finally, to start the procedure, only the initial coefficients and covariance matrix have to be constructed. There are additional settings required for exploration (included in the policy) and the adaptability (to be discussed in the next section). Note that there is an efficient method to derive the directional derivatives, which are used in the optimal policy, available in de Visser et al. [2011].

While the framework functions in an infinite-time setting, the simulation has a time limit after which a new trial is started. This to have a successful convergence to the optimal value function by visiting every state due to a random initial state $\mathbf{x}_0$. At each time step an action is selected, the next state is determined, the reward is calculated and the RLS TD algorithm updates the parameter coefficients $\mathbf{c}$ (Eq. 16) and covariance matrix $\mathbf{P}$ (Eq. 21). An overview of the complete algorithm is available in Table 3.

The requirement of each state $\mathbf{x} \in X$ being visited is essential to guarantee the convergence to the optimal coefficients $\mathbf{c}^*$. Therefore an explorer is introduced into the framework that explores the entire state-space X. Note that this proof assumes that the entire state-space is reachable, which is not true for every environment.

In this situation, the SDP framework consists of a direct implementation of the RLS TD algorithm, using multivariate simplex B-splines as a linear-in-the-parameters function approximator. As a consequence, the convergence proof as given in Bradtke and Barto [1996] applies.

A block diagram of the control scheme is presented in Figure 2. Both the state $\mathbf{x}$ and input $\mathbf{u}$ are in the diagram, as well as the reward r and the system disturbance w. The DP components are the Policy, Reward and Value, however where both the Policy and Reward components are identical for NDP and SDP, the Value component differs. The method to construct the value function in the

Table 3: The SDP overview

Step	Action
(0)	Initialization
(0a)	Create a spline space $S_d^r(\mathcal{T})$
(0b)	Set the parameter values $\mathbf{c}_1 = \mathbf{0}$
(0c)	Set the covariance matrix $\mathbf{P}_1 = \beta_1 \cdot (\mathbf{I} - \mathbf{H}^+\mathbf{H})$
(1)	Do for $n = 1, \ldots,$ trials,
(1a)	Set the initial state $\mathbf{x}_1$
(2)	Do for $t = 1, \ldots,$ end,
(2a)	$\mathbf{u}_t = f(\mathbf{x}_t, \mathbf{c}_t)$
(2b)	$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t)$
(2c)	$r_{t+1} = f(\mathbf{x}_{t+1}, \mathbf{u}_t)$
(2d)	Update $\mathbf{c}$ and $\mathbf{P}$ with the RLS TD algorithm

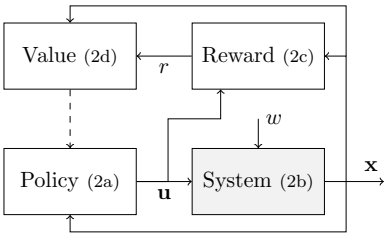

Figure 2: The control diagram of the SDP framework, including the corresponding step from Table 3

SDP framework is described in section 2.2.

4.2 Recursive weighted least squares

While the objective function in Eq. 4 will converge to the classic LS solution, it is unable to cope with time-varying systems as it weighs each measurement equally, driving the $\mathbf{P}$ matrix to zero. For the RLS algorithm to track time-varying systems, a popular and effective solution in adaptive control is using the forget factor, changing the quadratic objective function to:

$$\mathcal{J}_t = \frac{1}{t} \sum_{k=1}^{t} \beta^{t-k} \left[e_k\right]^2 \tag{23}$$

where β represents the forget factor. This equation can be rewritten as:

$$\mathcal{J}_t = \frac{1}{t} \left[\beta \mathcal{J}_{t-1} + \left[e_t\right]^2\right] \tag{24}$$

making it clear that that β has a discounting effect on the past errors, reducing the importance given to old data. Therefore, by applying a forget factor the LS solution is converted to a Weighted Least-Squares (WLS) solution, where the newest data-points have the most influence on parameters. According to Wellstead and Zarrop [1991] the forget factor can be applied to the covariance matrix as:

$$\mathbf{P}_{t+1} = \beta^{-1}\, \mathbf{P}_t \tag{25}$$

Applying the forget factor in this form has the disadvantage that it scales all elements of $\mathbf{P}$ equally. This will result in covariance wind-up when no new information is available over a long period, caused by certain elements of $\mathbf{P}$ becoming very large. This is a direct consequence of the spatial influence of the B-coefficients, visible in the B-net seen in Figure 1. The forget method perceived in Eq. 25 is therefore only used at initialization, represented by β_1 seen in Eq. 21.

A solution to prevent the covariance wind-up is to apply the forget factor only to the updated parameters. This approach is called directional forgetting Wellstead and Zarrop [1991] and updates the covariance matrix as follows:

$$\mathbf{P}_{t+1} = \mathbf{P}_t + \beta_2 \, \mathbf{B}_t \mathbf{B}_t^\top \tag{26}$$

where β_2 represents a forget factor applied to the updated parameters. However, the problem with this approach is that it destroys the continuity by ignoring the constraints set by $\mathbf{H}$ in Eq. 19. Therefore, in order to keep $\mathbf{c}$ in the null-space of $\mathbf{H}$, the following update is proposed:

$$\mathbf{P}_{t+1} = \mathbf{P}_t + \beta_2 \left[\mathbf{Z} \mathbf{B}_t \mathbf{B}_t^\top \mathbf{Z} \right] \tag{27}$$

where $\mathbf{Z} = (\mathbf{I} - \mathbf{H}^+ \mathbf{H})$, which is the projection on the null-space of $\mathbf{H}$, introduced in Eq. 21.

This approach can be implemented by altering step (2) in the original RLS algorithm from Table 2 to:

$$\mathbf{P}_{t+1} = \mathbf{P}_t - \frac{\mathbf{P}_t \mathbf{B}_t \mathbf{B}_t^\top \mathbf{P}_t}{1 + \mathbf{B}_t^\top \mathbf{P}_t \mathbf{B}_t} + \beta_2 \left[\mathbf{Z} \mathbf{B}_t \mathbf{B}_t^\top \mathbf{Z} \right] \tag{28}$$

resulting in our new formulation for the covariance matrix update step. The modified RLS TD algorithm with the additional term in step (2) is visible in Table 4, including the computational complexity. To immediately apply the forget factor at time t, step (3) employs the $\mathbf{P}_{t+1}$ matrix, as done in Ljung and Söderström [1983].

The RLS algorithm with directional forgetting is simply convergent for a system where the data generation mechanism is deterministic Bittanti et al. [1990]. It should be noted that under this assumption, LMS algorithms also have proven convergence Tsitsiklis and Van Roy [1997].

Additionally, the modified RLS TD algorithm is capable of filtering out the residual noise to end up near the optimal coefficients $\mathbf{c}^*$. With $\beta_2 = 0$, the filter has an infinite window in time, while at $\beta_2 > 0$, the window is infinite no longer, which has the advantage of being able to track time-varying systems and disadvantage of being susceptible to noise. This trade-off between noise filtering and tracking is an often returning phenomenon in adaptive control Wellstead and Zarrop [1991]. In principle $\beta_2 > 0$ only has a beneficial effect on the control of a time-varying system. Luckily, this approach allows the use of a variable β_2, able to increase and decrease as desired. The design of a successful variable forget factor will result in both good tracking behavior and a good performance with residual noise.

5 Performance Evaluation of SDP

The proposed SDP framework has been implemented on a pendulum swing-up non-linear control problem, as seen in A. At the start, the controller has

Table 4: RLS TD algorithm, modified from Table 1

Step		**Action**	**Computational Complexity**
(1)	$e_t =$	$r_{t+1} - (\mathbf{B}_t - \gamma \mathbf{B}_{t+1})^\top \mathbf{c}_t$	$\mathcal{O}(\hat{d})$
(2)	$\mathbf{P}_{t+1} =$	$\mathbf{P}_t - \frac{\mathbf{P}_t \mathbf{B}_t (\mathbf{B}_t - \gamma \mathbf{B}_{t+1})^\top \mathbf{P}_t}{1 + (\mathbf{B}_t - \gamma \mathbf{B}_{t+1})^\top \mathbf{P}_t \mathbf{B}_t}$ $+\beta_2 \left[\mathbf{Z} \mathbf{B}_t \mathbf{B}_t^\top \mathbf{Z} \right]$	$\mathcal{O}(\hat{a}^2)$
(3)	$\mathbf{c}_{t+1} =$	$\mathbf{c}_t + \mathbf{P}_{t+1} \mathbf{B}_t e_t$	$\mathcal{O}(\hat{a}^2)$

no knowledge about the optimal value function, and has to learn from online measurements. The gain input of the plant is assumed to be known; it will increase the learning time for both algorithms with the same amount if it is to be identified using model identification. The objective of the experiment is to move and keep the pendulum in an upwards position by using a limited torque. The controller receives reinforcement at each state, where the top position is most beneficial. The system dynamics are simulated using an Euler integration scheme in combination with the equations of motion as presented in A.

The performance of each trial is measured by the maximum amount of time t_{up} the pendulum is consecutively kept in an upwards position, where the upwards position is defined as:

$$|\theta_{up}| < \frac{\pi}{4} = 45^\circ \tag{29}$$

A trial itself consists of 20 s, with time steps of 0.02 s. As each trial is initialized in a random angle θ and a zero angle rate $\dot{\theta}$ (consistent with the experiment in Morimoto and Doya [2005]), some trials require more swings to reach the top. Therefore a lower t_{up} does not mean a worse performance per se, but it may have been initialized in a lower position. However, while θ_0 is random, it is identical for each trial over the different methods. This is done to remove the chance of one method having better initializations than the other, degrading the quality of the comparison.

The neural networks used for NDP are constructed using either radial basis function or using a sigmoid function as a basis. This corresponds to the radial basis network and feedforward network respectively. In case of the radial basis network, the TD(λ) algorithm is used for training to increase the performance, while the feedforward network is trained using the gradient descent approach. More information on how these networks are constructed and trained is available in Bertsekas and Tsitsiklis [1996] and Rojas [1996].

For SDP, a 4^{th} degree spline space with 1^{st} order continuity, without ($\beta_2 = 0$) and with ($\beta_2 = 0.4$) forget factor (see Eq. 28) has been selected. The polynomial degree has been determined by trial and error, such that the simplex spline is capable of estimating the optimal value function. Because the optimal policy is based on the first derivative and $\mathbf{u}$ has no rate restrictions, a discontinuous first derivative would give an unfair advantage to SDP. Therefore the continuity degree has been chosen such that the first derivative is continuous, identical to NDP. The value of β_2 is selected such that an increase of tracking behavior is witnessed in the experiments. A type III Delaunay triangulation of nodes positioned in a grid is used to produce the triangulation $\mathcal{T}_{32}$ seen in Figure 3.

Table 5: The Learning Parameters of the NDP and SDP Methods used to obtain the Results

NDP - Feedforward

Parameter	η_1	η_2	Neurons	Total parameters
Value	10^{-3}	10^{-3}	160	480

NDP - Radial basis

Parameter	η_1	η_2	λ	Neurons	Total parameters
Value	10^{-3}	10^{-2}	0.8	12x12	432

SDP

Parameter	S_d^r	β_1	β_2	Triangulation	Total parameters
Value	S_4^1	10	0 / 0.4	Type III $\mathcal{T}_{32}$	480

As explained in section 4.1, it is essential that this triangulation is capable of approximating the optimal value function.

The parameters of NDP used in the simulation have been selected such that there is a comparable amount of parameters in each function approximator; the $S_4^1(\mathcal{T}_{32})$ has a total of 480 coefficients, the feedforward network has 480 weights, and the 12x12 radial basis network has 432 weights. The centers of the radial basis network are positioned in a grid, including one centered in $\mathbf{x} = [0\ 0]^\top$. A result of the continuity constraints is that the number of free parameters is less than the total amount of parameters, thus effectively lowering the approximation power. In this case, there are 151 free parameters as the rank of $\mathbf{H}$ is 329.

A search for the best set of learning parameters for the feedforward and radial basis network was performed in an attempt to have a strong comparison between NDP and SDP. An overview of the parameters used in the experiments is visible in Table 5. The initialization of network weights, center weights or nodes has a significant impact; for the feedforward network and simplex splines half of the initializations failed when the networks weights or nodes were selected randomly, while for the radial basis network only 4% failed. In this case, success is described as scoring a $t_{up} > 10$ s for at least one trial. In the experiment the center weights and nodes were defined a-priori, removing the dependency on the initialization.

First, in section 5.1 the four control methods are tasked with controlling a stochastic system, which will demonstrate the influence of system noise. Secondly, in section 5.2 the methods are tasked with controlling a time-varying system; this is meant to exhibit the adaptability of the control methods.

5.1 Experiment I - The Stochastic System

In the first experiment the four methods are tasked with controlling a stochastic system. The stochastic system has a system disturbance of $\sigma_w = 3^\circ/s^2$. This means that the standard deviation of the system noise is 60% of the input's influence, as the maximum influence is $T^{\max}/ml^2 = 5^\circ/s^2$. Since the system noise is applied to the equation of motion, the effect of the noise is only directly connected to $\ddot{\theta}$. Furthermore, in order to identify the influence of system noise

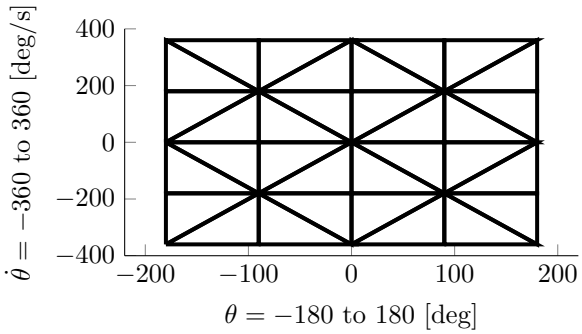

Figure 3: Type III Delaunay Triangulation $\mathcal{T}_{32}$

Table 6: Mean and standard deviation (std) of t_{up} of the 100 trials of Experiment I

$\sigma_w = 0°/s^2$

	NDP - FF	NDP - RBF	SDP - $\beta_2 = 0$	SDP - $\beta_2 = 0.4$
Mean	12.71 s	13.60 s	18.30 s	18.16 s
Std	7.89 s	7.80 s	2.84 s	3.24 s

$\sigma_w = 3°/s^2$

	NDP - FF	NDP - RBF	SDP - $\beta_2 = 0$	SDP - $\beta_2 = 0.4$
Mean	11.95 s	15.48 s	18.14 s	17.86 s
Std	6.62 s	4.93 s	3.08 s	2.91 s

on each method, the deterministic system ($\sigma_w = 0°/s^2$) is used as a baseline.

The results of the four methods are visible in the figures 4 and 6 for the deterministic system, and in the figures 5 and 7 for the stochastic system. It shows that while the learning parameters and initialized weights are identical for both systems, the NDP radial basis has learned to swing the pendulum up in less trials for the stochastic system. This increased learning rate of a dynamic programming algorithm in a stochastic system is a well known phenomenon, and is a result of the extra exploration that occurs due to the system noise Bertsekas and Tsitsiklis [1996]. Nevertheless the learning rate of SDP remains the highest in both the deterministic and stochastic system.

Another observation is that the stability of NDP is effected most by the presence of system noise, visible by an overall decrease of t_{up}. For SDP, only SDP - $\beta_2 = 0.4$ shows a decrease of t_{up} larger than 5 s in trial 36 and 68, thus it can be concluded that SDP shows most resilience to system noise. To support this claim with numerical arguments both the mean and standard deviation of t_{up} of the 100 trials are presented in Table 6. While it does not reveal the decrease in stability of the NDP methods because it is masked by the increased learning rate, it does show that SDP outperforms NDP in both scenarios.

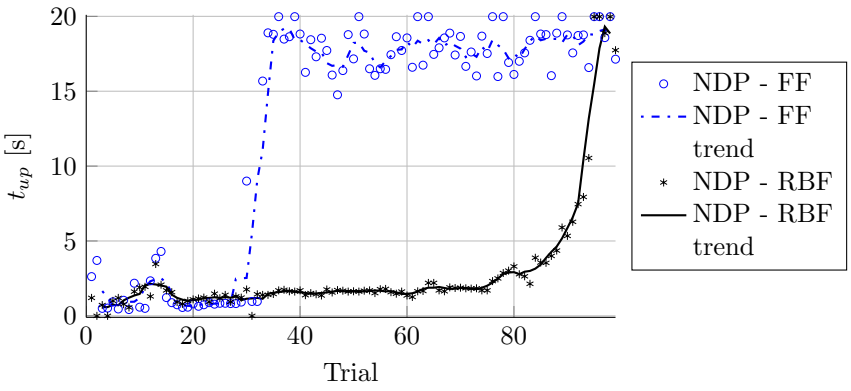

Figure 4: Performance overview of NDP - Radial basis 12x12 (432 parameters) and feedforward network 160 neurons (480 parameters), without system noise ($\sigma_w = 0°/s^2$). The trendline respresents a 5 point moving average.

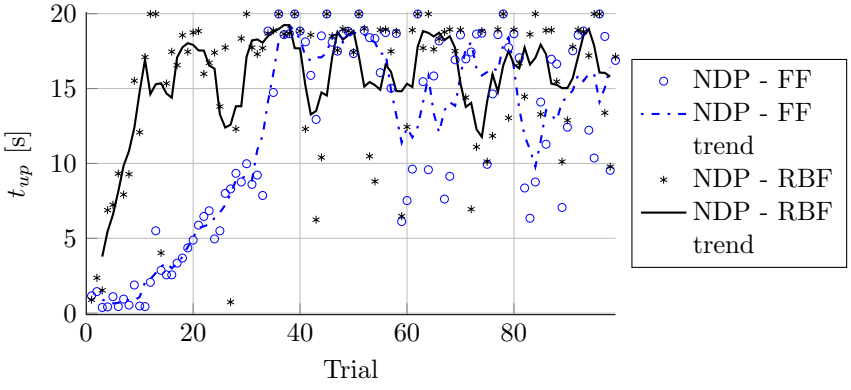

Figure 5: Performance overview of NDP - Radial basis 12x12 (432 parameters) and feedforward network 160 neurons (480 parameters), with system noise ($\sigma_w = 3°/s^2$). The trendline respresents a 5 point moving average.

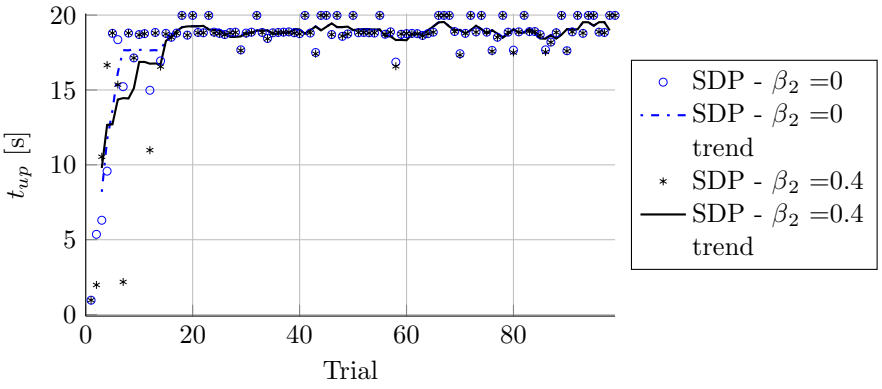

Figure 6: Performance overview of SDP - $S_4^1(\mathcal{T}_{32})$ (480 parameters), with ($\beta_2 = 0.4$) and without ($\beta_2 = 0$) forget factor, and without system noise ($\sigma_w = 0°/s^2$). The trendline respresents a 5 point moving average.

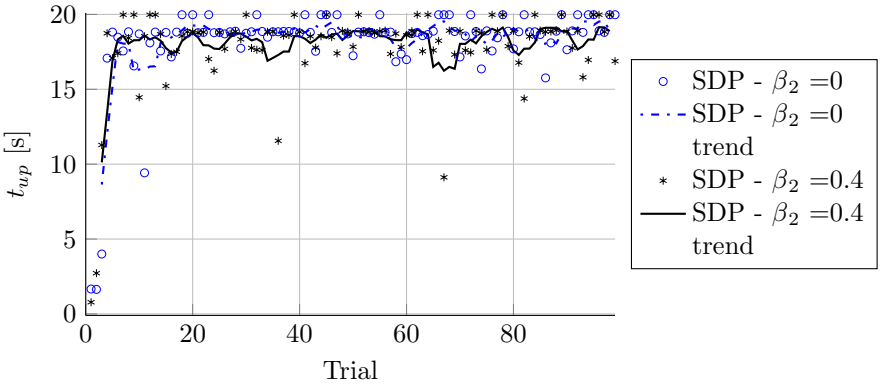

Figure 7: Performance overview of SDP - $S_4^1(\mathcal{T}_{32})$ (480 parameters), with ($\beta_2 = 0.4$) and without ($\beta_2 = 0$) forget factor, and with system noise ($\sigma_w = 3°/s^2$). The trendline respresents a 5 point moving average.

Table 7: Mean and standard deviation (std) of t_{up} of the 100 trials of Experiment II

	NDP - FF	NDP - RBF	SDP - $\beta_2 = 0$	SDP - $\beta_2 = 0.4$
Mean	17.88 s	11.44 s	12.40 s	15.90 s
Std	1.47 s	7.89 s	6.76 s	5.15 s

5.2 Experiment II - The Time-Varying System

The second experiment involves the control of a time-varying system. To simulate this, the control system has first been allowed to converge to the optimal value function by executing 1000 learning trials. The pendulum's mass is then changed from $m = 1$ kg to $m = 1.5$ kg and then 100 trials are simulated, similar to the first experiment.

The t_{up} of 100 trials after the change are visible in figures 8 and 9 for the NDP and SDP methods respectively. By increasing the pendulum's mass by 50% the old control system does not stop working, in fact, in many cases the top can still be reached, albeit not as close to $\theta = 0$ as before the change (i.e. the pendulum is held stationary at a slight angle). The changed optimal value function introduces a TD-error which propagates through the network. For the feedforward network a minor adaptation of the parameters is sufficient to adapt the global shape of the estimated value function, allowing the NDP - feedforward to continue controlling the altered system without temporary decreased performance. This is in contrast to NDP - radial basis, which does produce a decrease in performance as the TD-error propagates through the estimated value function. SDP - $\beta_2 = 0$ gives each data-point an equal weight, it is slow in adapting itself to a new situation, spending a long period in the transition phase where performance is reduced. SDP - $\beta_2 = 0.4$ has a shorter transition phase and has adapted itself to the new system before the 50^{th} trial.

In Table 7, the mean and standard deviation of t_{up} of the 100 trials are presented. This quantification identifies NDP - feedforward as the method with the best performance, and SDP - $\beta_2 = 0.4$ as the method with the second best performance. As stated before, the performance of the feedforward network is a direct consequence of its ability to generalize. However, there is no guarantee of convergence. Furthermore, it identifies NDP - radial basis and SDP - $\beta_2 = 0$ as equally bad, however in the figure 8 it can be seen that NDP - radial basis has recovered from the system change in the last 20 trials. This indicates that NDP - radial basis is capable of recovering although it takes more trials than the other methods.

6 Discussion

The most important difference between neural networks and simplex splines for a DP framework, is that neural networks are non-linear in the parameters, while simplex splines are linear-in-the-parameters. Using a linear-in-the-parameters function approximator allows for the use of the RLS TD algorithm, which has a fast and proven convergence in a stochastic framework Bradtke and Barto [1996]. As a result, the SDP framework without forget factor ($\beta_2 = 0$) has proven con-

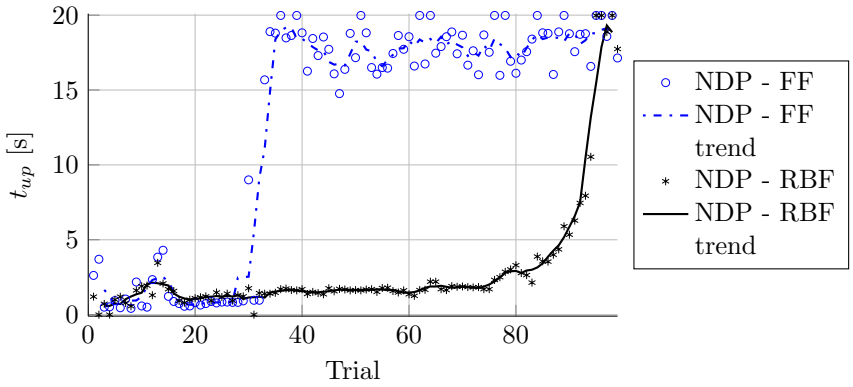

Figure 8: Performance overview of NDP - Radial basis 12x12 (432 parameters) and feedforward network 160 neurons (480 parameters), with an altered system parameter ($m = 1.5$) at $t = 0$ and without system noise ($\sigma_w = 0°/s^2$). The trendline respresents a 5 point moving average.

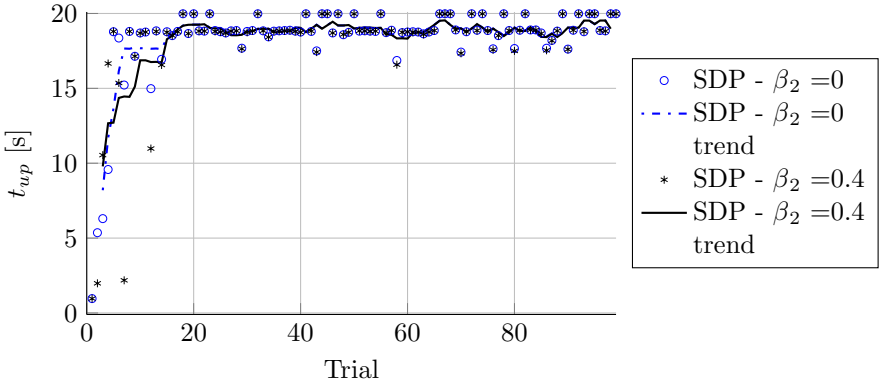

Figure 9: Performance overview of SDP - $S_4^1(\mathcal{T}_{32})$ (480 parameters), with ($\beta_2 = 0.4$) and without ($\beta_2 = 0$) forget factor, with an altered system parameter ($m = 1.5$) at $t = 0$ and without system noise ($\sigma_w = 0°/s^2$). The trendline respresents a 5 point moving average.

vergence demonstrates stable performance when learning online. Nevertheless, in certain circumstances it is beneficial to trade these properties for adaptability, which is done by using the modified RLS TD algorithm from Table 4 and thus introducing a forget factor ($\beta_2 > 0$). These circumstances arise when the environment is susceptible to system parameter changes and quick adaptation is required.

To successfully implement the SDP framework, it is important to construct the proper spline space. Because, even though SDP has proven convergence to the best fit, there is no guarantee of system performance. To obtain this guarantee, the optimal value function, or its shape, must be known a-priori. In practice this will mean that either an off-line simulation is used, or the system is tested to see if the desired performance is reached. Another option is to treat all unknown parameters as an additional optimization, and solve the entire optimization problem using Intersplines de Visser et al. [2012]. Unfortunately, at the moment Intersplines are limited to two-dimensional inputs, and require too much calculation power to make it attractive for real-time applications.

While it is possible to use the multivariate simplex B-splines in higher dimensions de Boor [1986], there are two problems that arise. First there is the construction of the triangulation, which is not automated and is already tedious in a dimension $N = 3$. Secondly, due to the "curse of dimensionality" Bellman [1957], the computational costs of dynamic programming are very high when moving to higher dimensions. The effects of this curse are reduced by selecting a triangulation that has fewer simplices, but retains the ability to represent the optimal value function. However, the first argument was that constructing a triangulation becomes increasingly difficult at higher dimensions and is not yet automated. This creates difficulties when applying the SDP framework to a problem with the dimension $N > 2$, since a similar grid-approach as used in the simulation, produces high computation costs.

A final remark on the methods used in the experiments. The four methods were considered to be comparable in terms of the number of parameters. However, the continuity constraints reduced the amount of free parameters in the SDP methods; thereby reducing the approximation power of the splines relative to the neural networks. Experiments with NDP have shown that reducing the amount of parameters reduces the performance, which implies that SDP is potentially better than concluded in the experiments.

7 Conclusion

In this paper the SDP framework was introduced; a combination of the RLS TD algorithm and multivariate simplex B-splines. It was shown that it is capable of solving the non-linear control problem of the pendulum swing-up with nothing but a reward function as feedback. This was done in significantly less trials than NDP systems supported by function approximators with a comparable amount of parameters. In addition, SDP presented a greater resilience to system noise than NDP, demonstrating no decrease in performance in the presence of a disturbance. Furthermore, a forget method that preserves the continuity constraints is introduced, which is merged with the RLS TD algorithm to create an adaptive control system. In conclusion it can be said that the high convergence rate of the RLS TD algorithm, in combination with the high

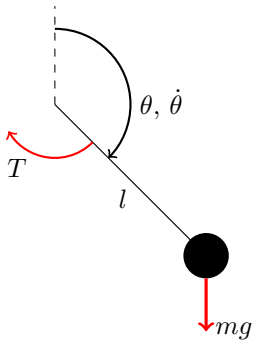

Figure 10: Pendulum Swing-Up

approximation power of the multivariate simplex B-splines, provides a basis for high performance non-linear control at the cost of a higher computational load requirement. By using the modified RLS TD algorithm, it is even possible to keep track of time-varying systems, resulting in an adaptive control method for non-linear systems.

In order to have a good trade-off between noise filtering and adaptability, it is important to design a forget factor that is capable of detecting a parameter change in the controlled system. There is extensive literature available on this subject and it is recommended to investigate which is the most effective in the SDP framework.

The unsolved issue of multivariate simplex B-splines remains the search for the optimal triangulation. While the assumed static triangulation performs adequately, as seen in the experiments, much can be gained in terms of computational complexity and performance by optimizing the triangulation. The advantage of finding the optimal triangluation grows exponentially, as it is connected to the "Curse of Dimensionality".

A Pendulum Swing-Up Task

The dynamics of the non-linear control task determined by:

$$ml^2\ddot{\theta} = -\mu\dot{\theta} + mgl\sin\theta + T \tag{30}$$

where θ is the angle from upright position, T is the limited input torque, $\mu = 0.01$ is the friction coefficient, $m = 1.0kg$ is the point mass at the tip of the rod, $l = 1.0m$ is the length of the pendulum, and $g = 9.8m/s^2$ is the gravity acceleration. A schematic overview is visible in Fig. 10.

The state vector is defined as $\mathbf{x} = [\theta\ \dot{\theta}]^\top$ and the action as $\mathbf{u} = T$. The complete system description is now:

$$\dot{\mathbf{x}} = f(\mathbf{x},\mathbf{u}) = \begin{pmatrix} \dot{\theta} \\ \frac{g}{l}\sin\theta - \frac{\mu}{ml^2}\dot{\theta} \end{pmatrix} + \begin{pmatrix} 0 \\ \frac{1}{ml^2} \end{pmatrix}\mathbf{u} + \begin{pmatrix} 0 \\ 1 \end{pmatrix} w \tag{31}$$

where w represents white noise with a standard deviation of σ_w. The reward

function of the system is:

$$r(\mathbf{x}, \mathbf{u}) = c_x(\cos\theta - 1) + c_u \int_0^{\frac{T}{T^{\max}}} \tan(\frac{\pi}{2}s)ds \tag{32}$$

where $c_x = 1$, $c_u = 0.1$ and $T^{\max} = 5$. As determined in section 2.3, the optimal control law is:

$$\mathbf{u} = T^{\max} \tanh(\frac{\pi}{2}\frac{1}{c_j}\tau\frac{\partial V^*}{\partial \mathbf{x}}\frac{\partial f(\mathbf{x}, \mathbf{u})}{\partial \mathbf{u}} + n) \tag{33}$$

where n represents white noise with a standard deviation of $\sigma_n = 0.01$, used as an explorer.

References

C. W. Anderson, P. M. Young, M. R. Buehner, J. N. Knight, K. A. Bush, and D. C. Hittle. Robust Reinforcement Learning Control Using Integral Quadratic Constrains for Recurrent Neural Networks. *IEEE Transaction on neural networks*, 18(4):993–1002, July 2007.

S. Bakin, M. Hegland, and M. Osborne. Parallel MARS algorithm based on B-splines. *Computational Statistics*, 15(4):463–484, 2000.

A. G. Barto, R. S. Sutton, and C. W. Anderson. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE transactions on systems, man and cybernetics*, 13(5):834–846, Sept. 1983.

R. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, 1957.

R. Bellman, R. Kalaba, and B. Kotkin. Polynomial Approximation - A New Computional Technique in Dynamic Programming: Allocation Processes. *Mathematics of Computation*, 17(82):155–161, Apr. 1963.

D. P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 2. Athena Scientific, 2007. ISBN 1-886529-30-2.

D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-dynamic programming*. Athena Scientific, 2 edition, 1996. ISBN 1-886529-10-8.

S. Bittanti, P. Bolzern, and M. Campi. Convergence and exponential convergence of identification algorithms with directional forgetting factor. *Automatica*, 26(5):929–932, 1990.

S. J. Bradtke and A. G. Barto. Linear Least-Squares Algorithms for Temporal Difference Learning. *Machine Learning*, 22:33–57, 1996.

L. Busoniu, R. Babuska, B. De Schutter, and D. Ernst. *Reinforcement Learning and Dynamic Programming Using Function Approximators*. CRC Press, 2010. ISBN 978-1-4398-2108-4.

C. Cervellera, A. Wen, and V. Chen. Neural network and regression spline value function approximations for stochastic dynamic programming. *Computers and Operations Research*, 34(1):70–90, 2007.

V. Chen. Application of orthogonal arrays and mars to inventory forecasting stochastic dynamic programs. *Computational Statistics and Data Analysis*, 30(3):317–341, 1999.

V. Chen, D. Ruppert, and C. Shoemaker. Applying experimental design and regression splines to high-dimensional continuous-state stochastic dynamic programming. *Operations Research*, 47(1):38–53, 1999.

J. W. Daniel. Splines and Efficiency in Dynamic Programming. *Mathematical Analysis and Applications*, 2(54):402–407, 1976.

C. de Boor. *B (asic)-spline basics*. Mathematics Research Center, University of Wisconsin-Madison, 1986.

C. C. de Visser and M. Verhaegen. Wavefront reconstruction in adaptive optics systems using nonlinear multivariate splines. *JOSA A*, 30(1):82–95, 2013.

C. C. de Visser, Q. P. Chu, and J. A. Mulder. A new approach to linear regression with multivariate splines. *Automatica*, 45(12):2903–2909, 2009.

C. C. de Visser, Q. P. Chu, and J. A. Mulder. Differential constraints for bounded recursive identification with multivariate splines. *Automatica*, 47: 2059–2066, 2011.

C. C. de Visser, E. van Kampen, Q. P. Chu, and J. A. Mulder. Intersplines: A New Approach to Globally Optimal Multivariate Splines Using Interval Analysis. *Reliable Computing*, 17(2):153–191, 2012.

K. Doya. Temporal Difference Learning in Continuous Time and Space. *Advances in Neural Information Processing Systems*, pages 1073–1079, 1996.

K. Doya. Reinforcement Learning in Continuous Time and Space. *Neural Computation*, 12:219–245, 2000.

S. Ferrari and R. F. Stengel. Online Adaptive Critic Flight Control. *Journal of Guidance, Control and Dynamics*, 27(5):777–786, 2004.

J. H. Friedman. Multivariate Adaptive Regression Splines. *The Annals of Statistics*, 19(1):1–67, 1991.

D. Han and S. N. Balakrishnan. Adaptive Critic Based Neural Networks for Control-Constrained Agile Missile Control. *Proceedings of the American Control Conference*, pages 2600–2604, 1999.

M. J. Lai and L. L. Schumaker. *Spline functions on triangulations*. Cambridge University Press, 2007. ISBN 978-0-521-87592-9.

C. L. Lawson and R. J. Hanson. *Solving least squares problems*. Prentice-Hall, 1974. ISBN 978-0898713565.

D. T. Lee and B. J. Schachter. Two algorithms for constructing a delaunay triangulation. *International Journal of Computer and Information Sciences*, 9(3):219–242, 1980.

L. Ljung and T. Söderström. *Theory and practice of recursive identification*. MIT Press, 1983. ISBN 13-978-0262120951.

J. Ma and W. B. Powell. A convergent recursive least squares approximate policy iteration algorithm for multi-dimensional markov decision process with continuous state and action spaces. *IEEE symposium on adaptive dynamic programming and reinforcement learning*, pages 66–73, 2009.

J. Morimoto and K. Doya. Robust Reinforcement Learning. *Neural Computation*, 2(17):335–359, 2005.

W. B. Powell. *Approximate Dynamic Programming*. Wiley, 2007. ISBN 978-0-470-17155-4.

R. Rojas. *Neural Networks*. Springer, 1996. ISBN 978-3540605058.

G. Saini and S. N. Balakrishnan. Adaptive Critic Based Neurocontroller for Autolanding of Aircraft. *Proceedings of the American Control conference*, pages 1081–1085, 1997.

P. J. Schweitzer. Generalized Polynomial Approximations in Markovian Decision Processes. *Journal of Mathematical Analysis and Applications*, 2(110): 568–582, 1985.

J. Si, A. Barto, W. Powell, and D. Wunsch. *Handbook of Learning and Approximate Dynamic Programming*. IEEE Press, 2004. ISBN 0-471-66054-0.

T. H. Summers, K. Kunz, N. Kariotoglou, M. Kamgarpour, S. Summers, and J. Lygeros. Approximate dynamic programming via sum of squares programming. In *European Control Conference (ECC)*, pages 191–197. IEEE, 2013.

R. S. Sutton. Learning to Predict by the Methods of Temporal Differences. *Machine Learning*, 3(1):9–44, 1988.

R. S. Sutton and A. G. Barto. *Reinforcement learning: an introduction*. The MIT Press, 1998. ISBN 978-0-262-19398-1.

J. N. Tsitsiklis and B. Van Roy. An Analysis of Temporal-Difference Learning with Function Approximation. *IEEE transactions on automatic control*, 42 (5):674–690, May 1997.

P. E. Wellstead and M. B. Zarrop. *Self-tuning Systems: Control and Signal Processing*. John Wiley & Sons, 1991. ISBN 0-471-92883-6.